

Computational Evidence that Self-regulation of Creativity is Good for Society

Liane Gabora (liane.gabora@ubc.ca)

University of British Columbia

Department of Psychology, Okanagan campus, Arts Building, 3333 University Way
Kelowna BC, V1V 1V7, CANADA

Simon Tseng (s.tseng@alumni.ubc.ca)

University of British Columbia

Department of Engineering, 5000-2332 Main Mall
Vancouver BC, V6T 1Z4, CANADA

Abstract

Excess individual creativity can be detrimental to society because creators invest in unproven ideas at the expense of propagating proven ones. Moreover, a proportion of individuals can benefit from creativity without being creative themselves by copying creators. We hypothesized that (1) societies increase their rate of cultural evolution by tempering the novelty-generating effects of creativity with the novelty-preserving effects of imitation, and (2) this is carried out by selectively rewarding and punishing creativity according to the value of the individuals' creative outputs. We tested this using an agent-based model of cultural evolution in which each agent self-regulated its invention-to-imitation ratio as a function of the fitness of its cultural outputs. In self-regulating societies, agents segregated into creators and imitators. The mean fitness of cultural outputs was higher than in non-self-regulating societies, and changes in diversity were rapid and more pronounced. We discuss limitations and possible social implications of our findings.

Keywords: Agent-based model; creativity; imitation; individual differences; self regulation; cultural evolution EVOC.

Introduction

It is commonly assumed that creativity is desirable, and the more creative one is, the better. Our capacity for self-expression, problem solving, and making aesthetically pleasing artifacts, all stem from our creative abilities. However, individuals often claim that their creativity is stifled by social norms, policies, and institutions. Moreover, our educational systems do not appear to prioritize the cultivation of creativity, and in some ways discourage it.

Perhaps there is an adaptive value to these seemingly mixed messages that society sends about the social desirability of creativity. Perhaps what is best for society is that individuals vary widely with respect to how creative they are, so as to ensure that the society as a whole both generates novel variants, and preserves the best of them. This paper provides a computational test of the hypothesis that society as a whole benefits when individuals can vary how creative they are in response to the perceived effectiveness of their ideas.

Definition and Key Features of Creativity

There are a plethora of definitions of creativity in the literature; nevertheless, it is commonly accepted that a core characteristic of creativity is the production of an idea or product that meets two criteria: originality or *novelty*, and appropriateness, adaptiveness, or *usefulness*, *i.e.*, relevance to the task

at hand (Guilford, 1950; Moran, 2011). Not only are humans individually creative, but we build on each other's ideas such that over centuries, art, science, and technology, as well as customs and folk knowledge, can be said to evolve. This cumulative building of new innovations on existing products is sometimes referred to as the ratchet effect (Tomasello, Kruger, & Ratner, 1993). Creativity has long been associated with personal fulfillment (May, 1975; Rogers, 1959), self-actualization (Maslow, 1959), and maintaining a competitive edge in the marketplace. Thus it is often assumed that more creativity is necessarily better.

However, there are significant drawbacks to creativity (Cropley, Cropley, Kaufman, & Runco, 2010; Ludwig, 1995). Generating creative ideas is difficult and time consuming, and a creative solution to one problem often generates other problems, or has unexpected negative side effects that may only become apparent after much effort has been invested. Creativity is correlated with rule bending, law breaking, and social unrest (Sternberg & Lubart, 1995; Sulloway, 1996), aggression (Tacher & Readdick, 2006), group conflict (Troyer & Younggreen, 2009), and dishonesty (Gino & Ariely, 2012). Creative individuals are more likely to be viewed as aloof, arrogant, competitive, hostile, independent, introverted, lacking in warmth, nonconformist, norm doubting, unconscientious, unfriendly (Batey & Furnham, 2006; Qian, Plucker, & Shen, 2010; Treffinger, Young, Selby, & Shepardson, 2002). They tend to be more emotionally unstable, and more prone to affective disorders such as depression and bipolar disorder, and have a higher incidence of schizophrenic tendencies, than other segments of the population (Andreason, 1987; Eysenck, 1993; Flaherty, 2005). They are also more prone to drug and alcohol abuse, as well as suicide (Jamison, 1993; Goodwin, 1998; Rothenberg, 1990; Kaufman, 2003). This suggests that there is a cost to creativity, both to the individual and to society.

Balancing Novelty with Continuity

Given the correlation between creativity and personality traits that are potentially socially disruptive, it is perhaps fortunate that in a group of interacting individuals, not all of them need be particularly creative for the benefits of creativity to be felt throughout the group. The rest can reap the rewards of the creator's ideas by copying them, buying from them, or sim-

ply admiring them. Few of us know how to build a computer, or write a symphony, but they are nonetheless ours to use and enjoy. Of course, if everyone relied on the strategy of imitating others rather than coming up with their own ideas, the generation of cultural novelty would grind to a halt. On the other hand, if everyone were as creative as the most creative amongst us, the frequency of the above-mentioned antisocial tendencies of creative people might be sufficiently high to interfere with cultural stability; *i.e.*, the perpetuation of cultural continuity. It is well known in theoretical biology that both novelty and continuity are essential for evolution, that is, for cumulative, open-ended, adaptive change over time.

This need for both novelty and continuity was demonstrated in an agent-based model of cultural evolution (Gabora, 1995). Novelty was injected into the artificial society through the invention of new actions, and continuity was preserved through the imitation of existing actions. When agents never invented, there was nothing to imitate, and there was no cultural evolution at all. If the ratio of invention to imitation was even marginally greater than 0, not only was cumulative cultural evolution possible, but eventually all agents converged on optimal cultural outputs. When all agents always invented and never imitated, the mean fitness of cultural outputs was also sub-optimal because fit ideas were not dispersing through society. The society as a whole performed optimally when the ratio of creating to imitating was approximately 2:1. Although results obtained with a simple computer model may have little bearing on complex human societies, the finding that extremely high levels of creativity can be detrimental to the society suggests that there may be an adaptive value to society's ambivalent attitude toward creativity.

This suggested that society as a whole might benefit from a distinction between the conventional workforce and what has been called a "creative class" (Florida, 2002) This was investigated in the model by introducing two types of agents: imitators, that only obtained new actions by imitating neighbors, and creators, that obtained new actions either by inventing or imitating (Gabora & Firouzi, 2012). It was possible to vary the probability that creators create versus imitate; thus whereas a given agent was either a creator or an imitator throughout the entire run, the proportion of creators innovating or imitating in a given iteration fluctuated stochastically. The mean fitness of ideas across the artificial society was highest when not all agents were creators. Specifically, there was a tradeoff between C , the proportion of creators to imitators in the society, and p , how creative the creators were). This provided further support for the hypothesis that society as a whole functions optimally when creativity is tempered with continuity.

We then hypothesized that society as a whole might perform even better if individuals are able to adjust how creative they are over time in accordance with their perceived creative success, through mechanisms such as selective ostracization of deviant behaviour unless accompanied by the generation

of valuable cultural novelty. In this way society might self-organize into a balanced mix of novelty generating creators and continuity perpetuating imitators, both of which are necessary for cumulative cultural evolution. A first step in investigating this hypothesis was to determine whether it is algorithmically possible to increase the mean fitness of ideas in a society by enabling them to self-regulate how creative they are.

The Computational Model

We investigated this using an agent-based model of cultural evolution referred to as "EVolution of Culture", abbreviated EVOC. It uses neural network based agents that (1) invent new ideas, (2) imitate actions implemented by neighbors, (3) evaluate ideas, and (4) implement successful ideas as actions. EVOC is an elaboration of Meme and Variations, or MAV (Gabora, 1995), the earliest computer program to model culture as an evolutionary process in its own right, as opposed to modeling the interplay of cultural and biological evolution¹. The approach was inspired by the genetic algorithm (Holland, 1975), or GA. The GA is a search technique that finds solutions to complex problems by generating a population of candidate solutions through processes akin to mutation and recombination, selecting the best, and repeating until a satisfactory solution is found. The goal behind MAV, and also behind EVOC, was to distil the underlying logic of not biological evolution but cultural evolution, *i.e.*, the process by which ideas adapt and build on one another in the minds of interacting individuals. Agents do not evolve in a biological sense—they neither die nor have offspring—but do in a cultural sense, by generating and sharing ideas for actions. In cultural evolution, the generation of novelty takes place through invention instead of through mutation and recombination as in biological evolution, and the differential replication of novelty takes place through imitation, instead of through reproduction with inheritance as in biological evolution. EVOC has been used to address such questions as how does the presence of leaders or barriers to the diffusion of ideas affect cultural evolution.

We now summarize briefly the architecture of EVOC in sufficient detail to explain our results; for further details we refer the reader to previous publications (Gabora, 1995; Gabora, Chia, & Firouzi, 2013; Gabora & Firouzi, 2012).

Agents

Agents consist of (1) a neural network, which encodes ideas for actions and detects trends in what constitutes a fit action, (2) a 'perceptual system', which observes and evaluates neighbours' actions, and (3) a body, consisting of six body parts which implement actions. The neural network is composed of six input nodes and six corresponding output nodes that represent concepts of body parts (LEFT ARM, RIGHT ARM, LEFT LEG, RIGHT LEG, HEAD, and HIPS),

¹The approach can thus be contrasted with computer models of how individual learning affects biological evolution (Best, 1999; Higgs, 1992; Hinton & Nowlan, 1992; Hutchins & Hazellhurst, 1991).

and seven hidden nodes that represent more abstract concepts (LEFT, RIGHT, ARM, LEG, SYMMETRY, OPPOSITE, and MOVEMENT). Input nodes and output nodes are connected to hidden nodes of which they are instances (*e.g.*, RIGHT ARM is connected to RIGHT.) Each body part can occupy one of three possible positions: a neutral or default position, and two other positions, which are referred to as active positions. Activation of any input node activates the MOVEMENT hidden node. Same-direction activation of symmetrical input nodes (*e.g.*, positive activation – which represents upward motion – of both arms) activates the SYMMETRY node.

Invention

An idea for a new action is a pattern consisting of six elements that dictate the placement of the six body parts. Agents generate new actions by modifying their initial action or an action that has been invented previously or acquired through imitation. During invention, the pattern of activation on the output nodes is fed back to the input nodes, and invention is biased according to the activations of the SYMMETRY and MOVEMENT hidden nodes. (Were this not the case there would be no benefit to using a neural network.) To invent a new idea, for each node of the idea currently represented on the input layer of the neural network, the agent makes a probabilistic decision as to whether the position of that body part will change, and if it does, the direction of change is stochastically biased according to the learning rate. If the new idea has a higher fitness than the currently implemented idea, the agent learns and implements the action specified by that idea. When “chaining” is turned on, an agent can keep adding new sub-actions and thereby execute a multi-step action, so long as the most recently-added sub-action is both an optimal sub-action and different from the previous sub-action of that action (Gabora et al., 2013).

Imitation

The process of finding a neighbour to imitate works through a form of lazy (non-greedy) search. The imitating agent randomly scans its neighbours, and adopts the first action that is fitter than the action it is currently implementing. If it does not find a neighbour that is executing a fitter action than its own current action, it continues to execute the current action.

Evaluation: The Fitness Function

Following Holland (1975), we refer to the success of an action in the artificial world as its *fitness*, with the caveat that unlike its usage in biology, here the term is unrelated to number of offspring (or ideas derived from a given idea). The fitness function used in these experiments rewards activity of all body parts except for the head, and symmetrical limb movement. Fitness of a single-step action F_n is determined as per Eq. 1. Total body movement, m , is calculated by adding the number of active body parts, *i.e.*, body parts not in the neutral position.

$$F_n = m + 1.5(s_a + s_t) + 2(1 - m_h) \quad (1)$$

$s_a = 1$ if arms move symmetrically; 0 otherwise

$s_t = 1$ if legs move symmetrically; 0 otherwise

$m_h = 1$ if head is stationary; 0 otherwise

Note that there are multiple optima. (For example an action can be optimal if either both arms move up or if both arms move down.) The fitness F_c of a multi-step action with n chained single-step actions (each with fitness F_n) is calculated by Eq. 2.

$$F_c = \sum_{k=1}^n \frac{F_n}{1.2^{n-1}} \quad (2)$$

Learning

Invention makes use of the ability to detect, learn, and respond adaptively to trends. Since no action acquired through imitation or invention is implemented unless it is fitter than the current action, new actions provide valuable information about what constitutes an effective idea. Knowledge acquired through the evaluation of actions is translated into educated guesses about what constitutes a successful action by updating the learning rate. For example, an agent may learn that more overall movement tends to be either beneficial (as with the fitness function used here) or detrimental, or that symmetrical movement tends to be either beneficial (as with the fitness function used here) or detrimental, and bias the generation of new actions accordingly.

The Artificial World

These experiments used a default artificial world: a toroidal lattice with 1024 cells each occupied by a single, stationary agent, and a von Neumann neighborhood structure. Creators and imitators were randomly dispersed.

A Typical Run

Fitness and diversity of actions are initially low because all agents are initially immobile, implementing the same action, with all body parts in the neutral position. Soon some agent invents an action that has a higher fitness than immobility, and this action gets imitated, so fitness increases. Fitness increases further as other ideas get invented, assessed, implemented as actions, and spread through imitation. The diversity of actions increases as agents explore the space of possible actions, and then decreases as agents hone in on the fittest actions. Thus, over successive rounds of invention and imitation, the agents’ actions improve. EVOC thereby models how “descent with modification” occurs in a purely cultural context.

Method

To test the hypothesis that the mean fitness of cultural outputs across society increases faster with social regulation (SR) than without it, we increased the relative frequency of invention for agents that generated superior ideas, and decreased

it for agents that generated inferior ideas. To implement this the computer code was modified as follows. Each iteration, for each agent, the fitness of its current action relative to the mean fitness of actions for all agents at the previous iteration was assessed. Thus we obtained the relative fitness (RF) of its cultural output. The agent's personal probability of creating, $p(C)$, was modified a function of RF as follows:

$$p(C)_n = p(C)_{n-1} \times RF_{n-1} \quad (3)$$

The probability of imitating, $p(I)$, was $1 - p(C)$. Thus when SR was on, if relative fitness was high the agent invented more, and if it was low the agent imitated more. $p(C)$ was initialized at 0.5 for both SR and non-SR societies. We compared runs with SR to runs without it.

Results

All data are averages across 250 runs. The mean fitness of the cultural outputs of societies with SR (the ability to self-regulate inventiveness as a function of inventive success) was higher than that of societies without SR, as shown in Figure 1. The fact that ideas got increasingly fitter over the course of a run was due to the fact that chaining was turned on; a fit action could always be made fitter by adding another sub-action. When chaining was turned off, societies with SR also outperformed without it, except that in both cases the mean fitness plateaued when all agents converged on optimally fit ideas (not shown).




Figure 1: This graph plots the mean fitness of implemented actions across all agents over the duration of the run with and without social regulation.

The typical pattern was observed with respect to the diversity, or number of different ideas: an increase as the space of possibilities is explored followed by a decrease as agents converge on fit actions. However, this pattern occurred earlier, and was more pronounced, in societies with SR than in societies without it, as shown in Figure 2. Note that, with chaining turned on, although the number of different actions decreases, the agents do not converge on a static set of ac-

tions; the set of implemented actions changes continuously as they find new, fitter actions.




Figure 2: This graph plots the diversity, or total number of different actions, across all agents over the duration of the run with and without social regulation.

Societies with SR ended up separating into two distinct groups: one that primarily invented, and one that primarily imitated, as illustrated in Figure 3. Thus the observed increase in fitness can indeed be attributed to increasingly pronounced individual differences in their degree of creative expression over the course of a run. Agents that generated superior cultural outputs had more opportunity to do so, while agents that generated inferior cultural outputs became more likely to propagate proven effective ideas rather than reinvent the wheel.

Discussion

Should society encourage creativity whenever possible, or with respect to creativity, can there be too much of a good thing? Are the needs of the individual for creative expression at odds with society's need to reinforce conventions and established protocols? EVOC agents are too rudimentary to suffer the affective penalties of creativity but the model incorporates another important drawback to creativity: time spent inventing is time not spent imitating. Creative agents, absorbed in their creative process, effectively rupture the fabric of the artificial society; they act as insulators that impede the diffusion of proven solutions. Imitators, in contrast, serve as a "cultural memory" that ensures the preservation of successful ideas. When effective inventors created more and poor inventors created less, the society as a whole could capitalize on the creative abilities of the best inventors and capitalize on efforts of the rest to disseminate fit cultural outputs. The results suggest that it can be beneficial for a social group if individuals are allowed to follow different developmental trajectories in accordance with their demonstrated successes.

These results do not prove that in real societies successful creators invent more and unsuccessful creators invent less; they merely show this kind of self-regulation is a feasible




Figure 3: At the beginning of the run (top) all agents created and imitated with equal probability. Midway through the run their $p(C)$ values were distributed along the range of values from 0 to 1. By the end of the run (bottom) they had segregated into imitators (with $p(C)$ from 0 to 0.1) and creators (with $p(C)$ from 0.9 to 1).

means of increasing the mean fitness of creative outputs. However, the fact that strong individual differences in creativity exist (Kaufman, 2003; Wolfradt & Pretz, 2001) suggests that this occurs in real societies. Whether prompted by individuals themselves or mediated by way of social cues, families, organizations, or societies may spontaneously self-organize to achieve a balance between creative processes that generate innovations and the imitative processes that disseminate these innovations. They evolve faster by tempering novelty with continuity. A more complex version of this scheme is that individuals find a task at which they excel and become unconventional with respect to that task only, such that for each task domain there exists some individual in the social group who comes to be best equipped to explore that space of possibilities.

The social practice of discouraging creativity until the individual has proven him- or herself may serve as a form of social self-regulation ensuring that creative efforts are not squandered. Individuals who are tuned to social norms and expectations, may over time become increasingly concerned with imitating and cooperating with others in a manner that promotes cultural continuity. Their thoughts travel more well-worn routes, and they are increasingly less likely to innovate. Others might be tuned to the demands of creative tasks, and less tethered to social norms and expectations, and thereby more likely to see things from unconventional perspectives. Thus they are more likely to come up with solutions to problems or unexpected challenges, find new avenues for self-expression, and contribute to the generation of cultural novelty. In other words, what Cropley *et al.* (2010) refer to as the “dark side of creativity” may reflect that the creative individual is tuned to task needs at expense of human needs. Although in the long run this benefits the group as a whole because it results in creative outputs, in the short run the creative individual may be less likely to obey social norms and live up to social expectations, and to experience stigmatization or discrimination as a result, particularly in his/her early years (Craft, 2005; Scott, 1999; Torrance, 1963). However, once the merits of such individuals’ creative efforts become known, they may be supported or even idolized.

Limitations of this work include that the fitness function was static throughout a run, and agents had only one action to optimize. In real life, there are many tasks, and a division of labor such that each agent specializes in a few tasks, and imitates other agents to carry out other tasks. It may be that no one individual is an across-the-board “creator” or “imitator” but that different individuals find different niches for domain-specific creative outputs.

Another limitation is that currently EVOC does not allow an agent to imitate some features of an idea and not others. This would be useful because cultural outputs both in EVOC and the real world exhibit a version of what in biology is referred to as epistasis, wherein what is optimal with respect to one component depends on what is going on with respect to another. Once both components have been optimized in a

mutually beneficial way (in EVOC, for example, symmetrical arm movement), excess creativity risks breaking up co-adapted partial solutions. In future studies we will investigate the effects of enabling partial imitation.

Acknowledgments

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada and the Flemish Fund for Scientific Research, Belgium.

References

- Andreason, N. (1987). Creativity and mental illness. prevalence rates in writers and their first degree relatives. *American Journal of Psychiatry*, *144*, 1288–1292.
- Batey, M., & Furnham, A. (2006). Creativity, intelligence, and personality: A critical review of the scattered literature. *Genetic and Social General Psychology Monographs*, *7*, 355–429.
- Best, M. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, *132*, 289–293.
- Craft, A. (2005). *Creativity in schools: Tensions and dilemmas*. London: Routledge.
- Cropley, D., Cropley, A., Kaufman, J., & Runco, M. (2010). *The dark side of creativity*. Cambridge UK: Cambridge University Press.
- Eysenck, H. (1993). Creativity and personality: Suggestions for a theory. *Psychological Inquiry*, *4*, 147–178.
- Flaherty, A. (2005). Frontotemporal and dopaminergic control of idea generation and creative drive. *Journal of Computational Neurology*, *493*, 147–153.
- Florida, R. (2002). *The rise of the creative class*. London: Basic Books.
- Gabora, L. (1995). Meme and variations: A computational model of cultural evolution. In L. Nadel & D. Stein (Eds.), *1993 lectures in complex systems*. Reading MA: Addison-Wesley.
- Gabora, L., Chia, W., & Firouzi, H. (2013). A computational model of two cognitive transitions underlying cultural evolution. In *Proceedings of the 35th annual meeting of the cognitive science society* (pp. 2344–2349). Houston TX: Cognitive Science Society.
- Gabora, L., & Firouzi, H. (2012). Society functions best with an intermediate level of creativity. In *Proceedings of the 34th annual meeting of the cognitive science society* (pp. 1578–1583). Houston TX: Cognitive Science Society.
- Gino, F., & Ariely, D. (2012). The dark side of creativity: Original thinkers can be more dishonest. *Journal of Personality and Social Psychology*, *102*, 445–459.
- Goodwin, D. (1998). *Alcohol and the writer*. New York: Penguin.
- Guilford, J. (1950). Creativity. *American Psychologist*, *5*, 444–454.
- Higgs, P. (1992). The mimetic transition: a simulation study of the evolution of learning by imitation. *Proceedings of the Royal Society B - Biological Sciences*, *267*, 1355–1361.
- Hinton, G., & Nowlan, S. (1992). How learning can guide evolution. *Complex Systems*, *267*, 495–502.
- Holland, J. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.
- Hutchins, E., & Hazelhurst, B. (1991). Learning in the cultural process. In C. Langton, J. Taylor, D. Farmer, & S. Rasmussen (Eds.), *Artificial life ii*. Redwood City: Addison-Wesley.
- Jamison, K. (1993). *Touched by fire: Manic-depressive illness and the artistic temperament*. New York: Free Press.
- Kaufman, J. (2003). The cost of the muse: Poets die young. *Death Studies*, *27*, 813–822.
- Ludwig, A. (1995). *The price of greatness*. New York: Guilford Press.
- Maslow, A. (1959). Creativity in self-actualizing people. In Harper & Brothers (Eds.), *Creativity and its cultivation*. New York: McGraw-Hill.
- May, R. (1975). *The courage to create*. New York: Bantam.
- Moran, S. (2011). The roles of creativity in society. In J. Kaufman & R. Sternberg (Eds.), *Cambridge handbook of creativity*. Cambridge UK: Cambridge University Press.
- Qian, M., Plucker, J., & Shen, J. (2010). A model of chinese adolescents creative personality. *Creativity Research Journal*, *22*, 62–67.
- Rogers, C. (1959). Toward a theory of creativity. In H. Anderson (Ed.), *Creativity and its cultivation*. New York: Harper & Row.
- Rothenberg, A. (1990). Creativity, mental health, and alcoholism. *Creativity Research Journal*, *3*, 179–201.
- Scott, C. (1999). Teachers biases toward creative children. *Creativity Research Journal*, *12*, 321–337.
- Sternberg, R., & Lubart, T. (1995). *Defying the crowd: Cultivating creativity in a culture of conformity*. New York: Free Press.
- Sulloway, F. (1996). *Born to rebel*. New York: Pantheon.
- Tacher, E., & Readdick, C. (2006). The relation between aggression and creativity among second graders. *Creativity Research Journal*, *18*, 261–267.
- Tomasello, M., Kruger, A., & Ratner, H. (1993). Cultural learning. *Behavioral and Brain Sciences*, *16*, 495–552.
- Torrance, E. (1963). *Guiding creative talent*. Englewood Cliffs, NJ: Prentice-Hall.
- Treffinger, D., Young, G., Selby, E., & Shepardson, C. (2002). *Assessing creativity: A guide for educators (rm02170)*. Storrs CT: University of Connecticut Press and The National Research Center on the Gifted and Talented.
- Troyer, L., & Younggreen, R. (2009). Conflict and creativity in groups. *Journal of Social Issues*, *65*, 409–413.
- Wolfradt, U., & Pretz, J. (2001). Individual differences in creativity: Personality, story writing, and hobbies. *European Journal of Personality*, *15*, 297–310.